Exhibit 14

Illustrative Claim Chart for U.S. Patent No. 10,628,273

Claim 1	Exemplary Dell Servers
A node system comprising:	The Exemplary Dell Servers are a node system. *See, e.g.,*:  "Dell PowerEdge Rack Servers."[1]
[a] a first computer that executes processing when operating as an active system of a redundant system;	The Exemplary Dell Servers comprise a first computer that executes processing when operating as an active system of a redundant system. For example, the Exemplary Dell '273 Products support High Availability (HA) and other redundant features that allow a first computer that executes processing when operating as an active system of a redundant system. *See, e.g.,*:

[1] Available at https://www.dell.com/en-us/dt/servers/poweredge-rack-servers.htm?hve=explore+poweredge-rack-servers#tab0=0&tab1=0&accordion0&accordion1&accordion2.

Claim 1	Exemplary Dell Servers

VMware vSphere High Availability

VMware vSphere High Availability delivers the availability required by most applications running in virtual machines, independent of the operating system and applications running in it. High Availability provides uniform, cost-effective failover protection against hardware and operating system outages within your virtualized IT environment. High Availability allows you to:

- Monitor VMware vSphere hosts and virtual machines to detect hardware and guest operating system failures.
- Restart virtual machines on other vSphere hosts in the cluster without manual intervention when a server outage is detected.
- Reduce application downtime by automatically restarting virtual machines upon detection of an operating system failure.

"VMware vSphere High Availability."[2]

Reliability

When an outage occurs, the last thing you want to worry about is whether a high availability solution will work. Guided by real-world customer feedback, VMware has added capabilities to maximize your confidence in High Availability, including the following:

- Elimination of external component dependencies—High Availability does not depend on DNS resolution. This reduces the likelihood that an external component outage will disrupt High Availability operations.
- Multiple communication paths—High Availability nodes within a cluster can communicate through the storage subsystem as well as over the management network. Multiple communication paths increase redundancy and enable better assessment of the health of a vSphere host and its virtual machines.
- VM-VM anti-affinity rules—High Availability respects VM-VM anti-affinity rules defined in VMware vSphere Distributed Resource Scheduler, eliminating the need for VMware vSphere vMotion migrations after failover.

[2] Available at https://www.vmware.com/products/vsphere/high-availability.html.

Claim 1	Exemplary Dell Servers
	Id. vSphere HA provides high availability for virtual machines by pooling the virtual machines and the hosts they reside on into a cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts. When you create a vSphere HA cluster, a single host is automatically elected as the primary host. The primary host communicates with vCenter Server and monitors the state of all protected virtual machines and of the secondary hosts. Different types of host failures are possible, and the primary host must detect and appropriately deal with the failure. The primary host must distinguish between a failed host and one that is in a network partition or that has become network isolated. The primary host uses network and datastore heartbeating to determine the type of failure. "How vSphere HA Works."[3]

[3] Available at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-33A65FF7-DA22-4DC5-8B18-5A7F97CCA536.html.

Claim 1	Exemplary Dell Servers
	> Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2 A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service. Failover Clustering has many practical applications, including: • Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines • Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V "Failover Clustering Overview."[4]

[4] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/failover-clustering-overview.

Claim 1	Exemplary Dell Servers
	 "What's New in Failover Clustering."[5]

[5] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/whats-new-in-failover-clustering.

Claim 1	Exemplary Dell Servers
	Windows Server Failover Clustering provides high availability for workloads running on Azure Stack HCI and Windows Server clusters. These resources are considered highly available if the nodes that host resources are up; however, the cluster generally requires more than half the nodes to be running, which is known as having *quorum*. Quorum is designed to prevent *split-brain* scenarios that can happen when there's a partition in the network and subsets of nodes can't communicate with each other. This can cause both subsets of nodes to try to own the workload and write to the same disk, which can lead to numerous problems. However, this is prevented with Failover Clustering's concept of quorum, which forces only one of these groups of nodes to continue running, so only one of these groups stays online. Quorum determines the number of failures that the cluster can sustain while still remaining online. Quorum is designed to handle the scenario when there's a problem with communication between subsets of cluster nodes, so that multiple servers don't try to simultaneously host a resource group and write to the same disk at the same time. By having this concept of quorum, the cluster forces the cluster service to stop in one of the subsets of nodes to ensure that there's only one true owner of a particular resource group. Nodes that have been stopped can once again communicate with the main group of nodes and will automatically rejoin the cluster and start their cluster service. In Azure Stack HCI and Windows Server 2019, there are two components of the system that have their own quorum mechanisms: • **Cluster Quorum**: This operates at the cluster level (i.e. you can lose nodes and have the cluster stay up) • **Pool Quorum**: This operates on the pool level (i.e. you can lose nodes and drives and have the pool stay up). Storage pools were designed to be used in both clustered and non-clustered scenarios, which is why they have a different quorum mechanism. "Understanding cluster and pool quorum."[6]

[6] Available at https://learn.microsoft.com/en-us/azure-stack/hci/concepts/quorum.

Claim 1	Exemplary Dell Servers
	# Fault domain awareness Article • 02/16/2023 > Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2 Failover Clustering enables multiple servers to work together to provide high availability – or put another way, to provide node fault tolerance. But today's businesses demand ever-greater availability from their infrastructure. To achieve cloud-like uptime, even highly unlikely occurrences such as chassis failures, rack outages, or natural disasters must be protected against. That's why Failover Clustering in Windows Server 2016 introduced chassis, rack, and site fault tolerance as well. "Fault domain awareness."[7]

[7] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/fault-domains.

Claim 1	Exemplary Dell Servers
	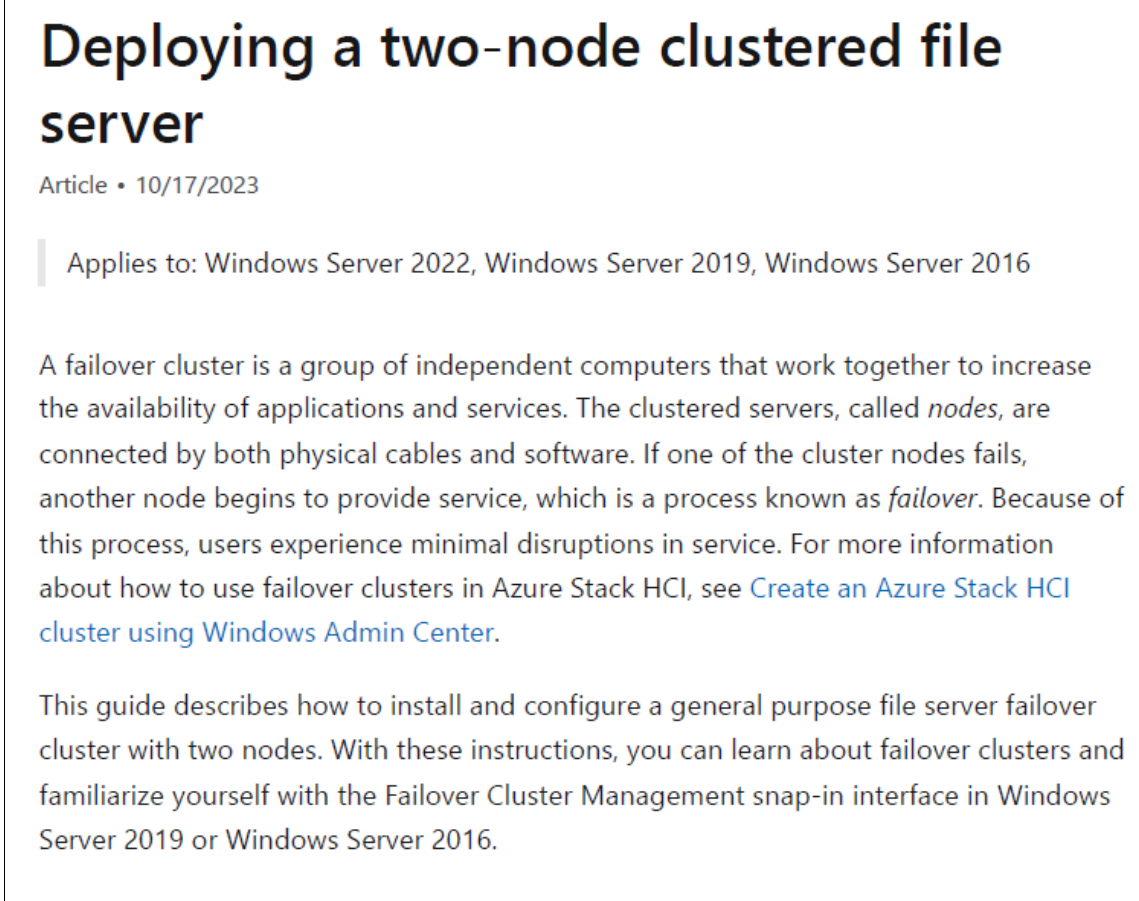

Deploying a two-node clustered file server

Article • 10/17/2023

> Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016

A failover cluster is a group of independent computers that work together to increase the availability of applications and services. The clustered servers, called *nodes*, are connected by both physical cables and software. If one of the cluster nodes fails, another node begins to provide service, which is a process known as *failover*. Because of this process, users experience minimal disruptions in service. For more information about how to use failover clusters in Azure Stack HCI, see Create an Azure Stack HCI cluster using Windows Admin Center.

This guide describes how to install and configure a general purpose file server failover cluster with two nodes. With these instructions, you can learn about failover clusters and familiarize yourself with the Failover Cluster Management snap-in interface in Windows Server 2019 or Windows Server 2016.

"Deploying a two-node clustered file server."[8]

[8] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/deploy-two-node-clustered-file-server?tabs=server-manager.

Claim 1	Exemplary Dell Servers
	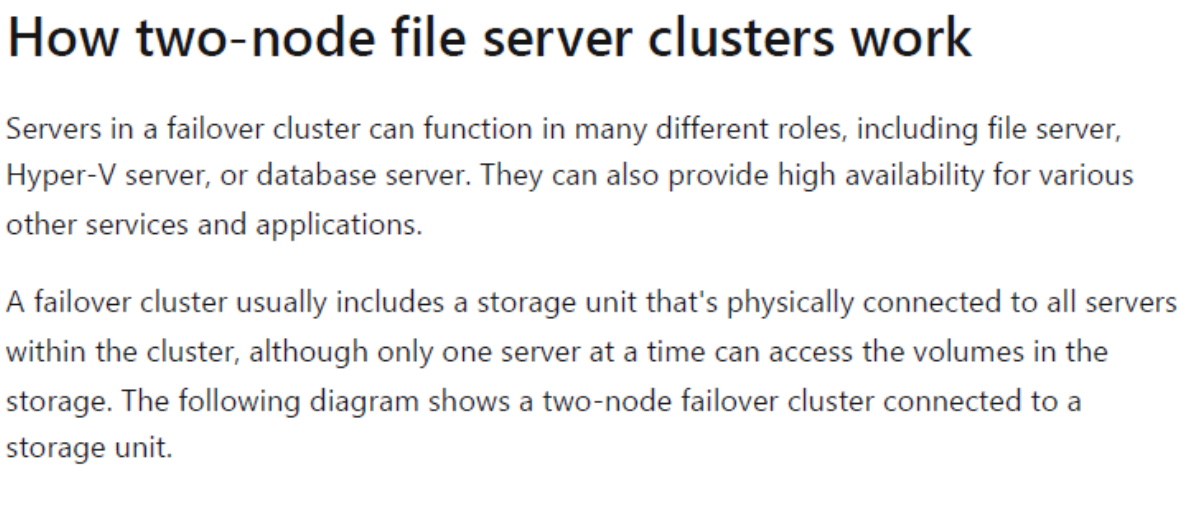 "How two-node file server clusters work."[9]

[9] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/deploy-two-node-clustered-file-server?tabs=server-manager.

Claim 1	Exemplary Dell Servers

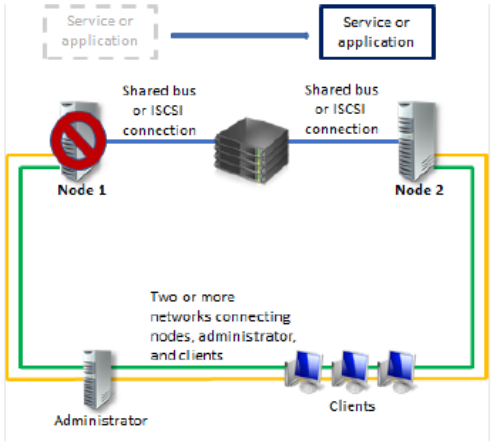

Storage volumes or logical unit numbers (LUNs) exposed to the nodes in a cluster must not be exposed to other servers, including servers in another cluster. The following diagram shows an example of what isolated LUNs look like in a deployment.

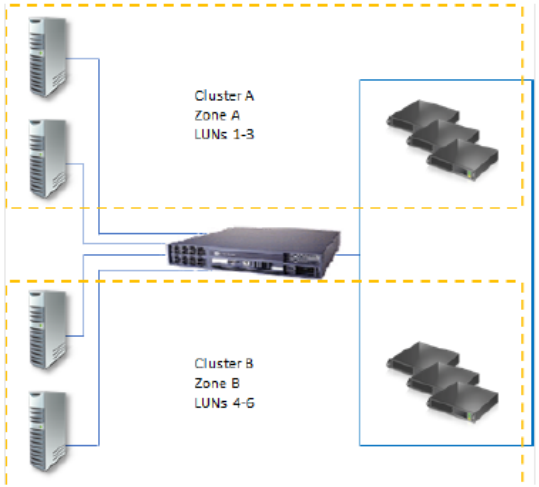

"Deploying a two-node clustered file server."

Claim 1	Exemplary Dell Servers
	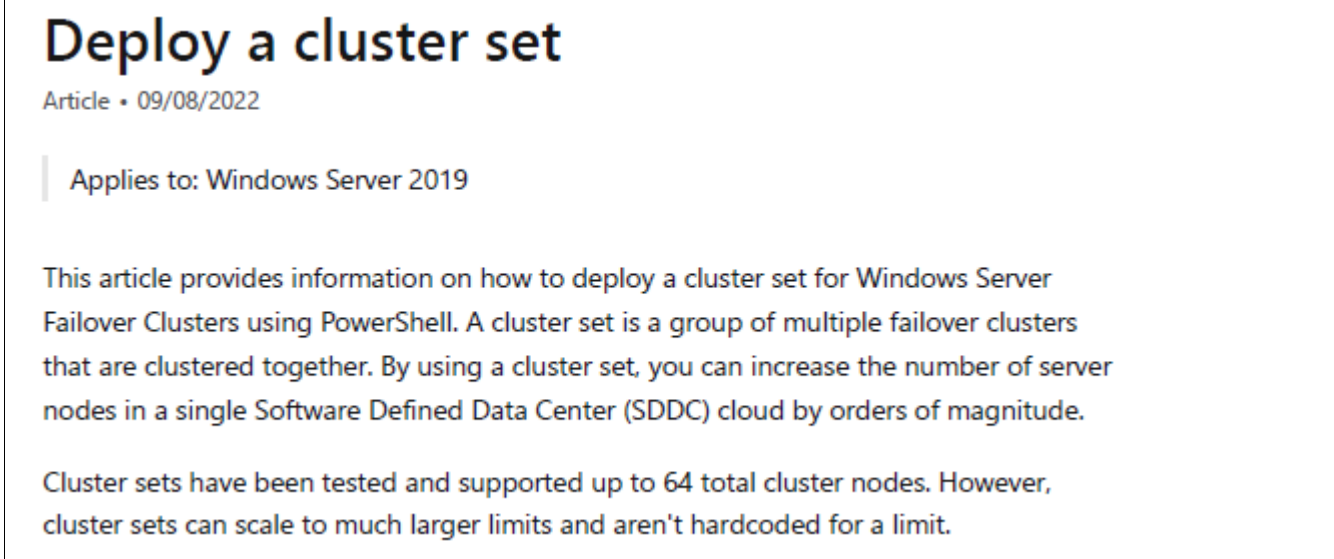 "Deploy a cluster set."[10]

[10] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/cluster-set.

Claim 1	Exemplary Dell Servers

Benefits

Cluster sets offer the following benefits:

- Significantly increases the supported SDDC cloud scale for running highly available virtual machines (VMs) by combining multiple smaller clusters into a single large fabric, while keeping the software fault boundary to a single cluster. You can easily migrate VMs across the cluster set.
- Increased resiliency. Having four 4-node clusters in a cluster set gives you better resiliency than a single 16-node cluster in that multiple compute nodes can go down and production remains intact.
- Management of failover cluster lifecycle, including onboarding and retiring clusters, without impacting tenant VM availability.
- VM flexibility across individual clusters and a present a unified storage namespace.
- Easily change the compute-to-storage workload ratio in your hyper-converged environment.
- Benefit from Azure-like Fault Domains and Availability sets across individual clusters in initial VM placement and subsequent migration.
- Can use even if compute and storage hardware between cluster nodes isn't identical.
- Live migration of VMs between clusters.
- Azure-like availability sets and fault domains across multiple clusters.

"Deploy a cluster set."

Claim 1	Exemplary Dell Servers

Use Cluster Shared Volumes in a failover cluster

Article • 02/11/2022

> Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Windows Server 2012, Windows Server 2012 R2, Azure Stack HCI, versions 21H2 and 20H2

Cluster Shared Volumes (CSV) enable multiple nodes in a Windows Server failover cluster or Azure Stack HCI to simultaneously have read-write access to the same LUN (disk) that is provisioned as an NTFS volume. The disk can be provisioned as Resilient File System (ReFS); however, the CSV drive will be in redirected mode meaning write access will be sent to the coordinator node. For more information, see About I/O synchronization and I/O redirection in CSV communication later in this document. With CSV, clustered roles can fail over quickly from one node to another node without requiring a change in drive ownership, or dismounting and remounting a volume. CSV also help simplify the management of a potentially large number of LUNs in a failover cluster.

CSV provides a general-purpose, clustered file system which is layered above NTFS or ReFS. CSV applications include:

- Clustered virtual hard disk (VHD/VHDX) files for clustered Hyper-V virtual machines
- Scale-out file shares to store application data for the Scale-Out File Server clustered role. Examples of the application data for this role include Hyper-V virtual machine files and Microsoft SQL Server data. Be aware that ReFS is not supported for a Scale-Out File Server in Windows Server 2012 R2 and below. For more information about Scale-Out File Server, see Scale-Out File Server for Application Data.
- Microsoft SQL Server 2014 (or higher) Failover Cluster Instance (FCI). Microsoft SQL Server clustered workload in SQL Server 2012 and earlier versions of SQL Server do not support the use of CSV.
- Windows Server 2019 or higher Microsoft Distributed Transaction Control (MSDTC)

"Use Cluster Shared Volumes in a failover cluster."[11]

[11] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/failover-cluster-csvs.

Claim 1	Exemplary Dell Servers
	# Review requirements and considerations for using CSV in a failover cluster Before using CSV in a failover cluster, review the network, storage, and other requirements and considerations in this section. ## Network configuration considerations Consider the following when you configure the networks that support CSV. - **Multiple networks and multiple network adapters**. To enable fault tolerance in the event of a network failure, we recommend that multiple cluster networks carry CSV traffic or that you configure teamed network adapters. "Use Cluster Shared Volumes in a failover cluster."

Claim 1	Exemplary Dell Servers

High Availability Clusters

High-availability clusters (also known as **HA clusters** , **failover clusters** or **Metroclusters Active/Active**) are groups of computers that support server applications that can be reliably utilised with a minimum amount of down-time.

They operate by using high availability software to harness redundant computers in groups or clusters that provide continued service when system components fail.

Without clustering, if a server running a particular application crashes, the application will be unavailable until the crashed server is fixed. HA clustering remedies this situation by detecting hardware/software faults, and immediately restarting the application on another system without requiring administrative intervention, a process known as failover.

As part of this process, clustering software may configure the node before starting the application on it. For example, appropriate file systems may need to be imported and mounted, network hardware may have to be configured, and some supporting applications may need to be running as well.

HA clusters are often used for critical databases, file sharing on a network, business applications, and customer services such as electronic commerce websites.

"Ubuntu – Introduction to High Availability."[12]

[12] Available at https://ubuntu.com/server/docs/introduction-to-high-availability.

Claim 1	Exemplary Dell Servers

High Availability cluster heartbeat

> HA cluster implementations attempt to build redundancy into a cluster to eliminate single points of failure, including multiple network connections and data storage which is redundantly connected via storage area networks.
>
> HA clusters usually use a heartbeat private network connection which is used to monitor the health and status of each node in the cluster. One subtle but serious condition all clustering software must be able to handle is split-brain, which occurs when all of the private links go down simultaneously, but the cluster nodes are still running.
>
> If that happens, each node in the cluster may mistakenly decide that every other node has gone down and attempt to start services that other nodes are still running. Having duplicate instances of services may cause data corruption on the shared storage.

"Ubuntu – Introduction to High Availability."

High Availability Cluster Quorum

> HA clusters often also use quorum witness storage (local or cloud) to avoid this scenario. A witness device cannot be shared between two halves of a split cluster, so in the event that all cluster members cannot communicate with each other (e.g., failed heartbeat), if a member cannot access the witness, it cannot become active.

"Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Example

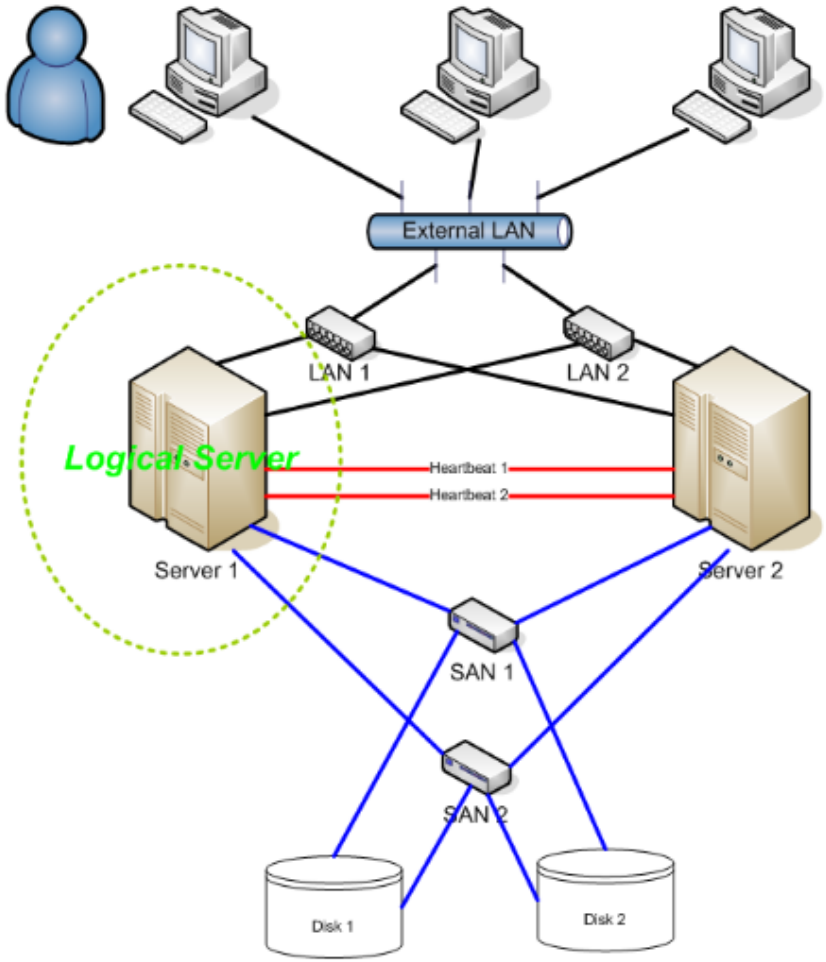

"Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Linux High Availability Projects

There are many upstream high availability related projects that are included in Ubuntu Linux. This section will describe the most important ones.

The following packages are present in latest Ubuntu LTS release:

Ubuntu HA Core Packages

Packages in this list are supported just like any other package available in **[main] repository** would be.

PACKAGE	URL
libqb	Ubuntu | Upstream
kronosnet	Ubuntu | Upstream
corosync	Ubuntu | Upstream
pacemaker	Ubuntu | Upstream
resource-agents	Ubuntu | Upstream
fence-agents	Ubuntu | Upstream
crmsh	Ubuntu | Upstream
pcs*	Ubuntu | Upstream
cluster-glue	Ubuntu | Upstream
drbd-utils	Ubuntu | Upstream
dlm	Ubuntu | Upstream
gfs2-utils	Ubuntu | Upstream
keepalived	Ubuntu | Upstream

"Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers
	• `libqb` - Library which provides a set of high performance client-server reusable features. It offers high performance logging, tracing, IPC and poll. Its initial features were spun off the `Corosync` cluster communication suite to make them accessible for other projects. • `Kronosnet` - `Kronosnet`, often referred to as `knet`, is a network abstraction layer designed for High Availability. `Corosync` uses `Kronosnet` to provide multiple networks for its interconnect (replacing the old Totem Redundant Ring Protocol) and add support for some more features like interconnect network hot-plug. • `Corosync` - or *Cluster Membership Layer*, provides reliable messaging, membership and quorum information about the cluster. Currently, Pacemaker supports `Corosync` as this layer. • **Pacemaker** - or *Cluster Resource Manager*, provides the brain that processes and reacts to events that occur in the cluster. Events might be: nodes joining or leaving the cluster, resource events caused by failures, maintenance, or scheduled activities. To achieve the desired availability, Pacemaker may start and stop resources and fence nodes. • **Resource Agents** - Scripts or operating system components that start, stop or monitor resources, given a set of resource parameters. These provide a uniform interface between pacemaker and the managed services. • **Fence Agents** - Scripts that execute node fencing actions, given a target and fence device parameters. • **crmsh** - Advanced command-line interface for High-Availability cluster management in GNU/Linux. • **pcs** - Pacemaker command line interface and GUI. It permits users to easily view, modify and create pacemaker based clusters. `pcs` also provides `pcsd`, which operates as a GUI and remote server for `pcs`. Together `pcs` and `pcsd` form the recommended configuration tool for use with pacemaker. *NOTE: It was added to the [main] repository in Ubuntu Lunar Lobster (23.10).* "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers
	• **cluster-glue** - Reusable cluster components for Linux HA. This package contains node fencing plugins, an error reporting utility, and other reusable cluster components from the Linux HA project. • **DRBD** - Distributed Replicated Block Device, **DRBD** is a distributed replicated storage system for the Linuxplatform. It is implemented as a kernel driver, several userspace management applications, and some shell scripts. DRBD is traditionally used in high availability (HA) clusters. • **DLM** - A distributed lock manager (DLM) runs in every machine in a cluster, with an identical copy of a cluster-wide lock database. In this way DLM provides software applications which are distributed across a cluster on multiple machines with a means to synchronize their accesses to shared resources. • **gfs2-utils** - Global File System 2 - filesystem tools. The Global File System allows a cluster of machines to concurrently access shared storage hardware like SANs or iSCSI and network block devices. • **Keepalived** - Keepalived provides simple and robust facilities for loadbalancing and high-availability to Linux system and Linux based infrastructures. Loadbalancing framework relies on well-known and widely used Linux Virtual Server (IPVS) kernel module providing Layer4 loadbalancing. Keepalived implements a set of checkers to dynamically and adaptively maintain and manage loadbalanced server pool according their health. On the other hand high-availability is achieved by VRRP protocol. "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Ubuntu HA Community Packages

Packages in this list are supported just like any other package available in **[universe] repository** would be.

PACKAGE	URL
pcs*	Ubuntu | Upstream
csync2	Ubuntu | Upstream
corosync-qdevice	Ubuntu | Upstream
fence-virt	Ubuntu | Upstream
sbd	Ubuntu | Upstream
booth	Ubuntu | Upstream

- **Corosync-Qdevice** - Its primary use is for even-node clusters, operates at corosync (quorum) layer. Corosync-Qdevice is an independent arbiter for solving split-brain situations. (qdevice-net supports multiple algorithms).

- **SBD** - A Fencing Block Device can be particularly useful in environments where traditional fencing mechanisms are not possible. SBD integrates with Pacemaker, a watchdog device and shared storage to arrange for nodes to reliably self-terminate when fencing is required.

> Note: **pcs** was added to the [main] repository in Ubuntu Lunar Lobster (23.04).

"Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Linux Clustering Concepts

A cluster is a group of computers (nodes) which work together to provide a shared solution. At a high level, a cluster can be viewed as having three parts (often defined as cluster stack).

Basic concepts

- **Resources:** These are the reason for the cluster's being the services that need to be kept highly available.
- **Resource Agents:** These are scripts operating system components that start, stop, and monitor resources, given a set of resource parameters.
- **Fence Agents:** These are scripts that execute a node fencing actions, given a target and fence device.
- **Fencing:** The ability to disable nodes.
- **Quorum:** Encapsulates the ability to determine whether the cluster can continue to operate safely or not.

Cluster types

The four types are as follows:

- **High Availability (HA):** Used for Fault Tolerance to keep server services available to employees or customers.
- **Load Balancing:** Balances the load between multiple systems when a service needs to be available to numerous systems at once (can be used for other three types of Clusters).
- **Distributed:** Jobs will be managed by different systems.
- **Parallel (Beowulf):** Jobs are managed by multiple processors on multiple systems.

As well by the configuration there are several types of clusters:

- **Manual clustering:** Lets you classify, merge, split clusters manually if the output of the automatic spike sorting algorithms are not satisfactory.
- **Merging clusters:** When multiple clusters seem to correspond to the same unit.
- **Splitting clusters:** You can create a new cluster by drawing a polygon around a set of spikes in the feature view, the amplitude view, the template amplitude view, or the spike attribute views.

"Rackspace – Clustering concepts."[13]

[13] Available at https://docs.rackspace.com/docs/clustering-concepts.

Claim 1	Exemplary Dell Servers
	Reasons to have a cluster There are different reasons for using a cluster. We use them to provide a resilient, highly available backend for solutions. This is primarily for backend services like MySQL, NFS or Redis where the service can be behind a front end of web servers. High availability is provided via automatic failover - if there is a failure on a cluster node, the clustered services running on that node will automatically relocate away to a node that is running properly. Clusters will generally use shared storage (SAN) so that data will be persistent when services move between nodes. In the event if a node becomes unresponsive, typically that node will be rebooted (fenced) by other nodes in order to preserve data integrity of the shared storage and floating IP ownership. **Linux Open Source High Availability Clustering** Some Linux operating system vendors offer clustering software, such as SUSE Linux HAE; Red Hat Enterprise Linux (RHEL); and Oracle Real Application Clusters (RAC). While they allow you to create a failover cluster, they present a variety of challenges, thist is highly manual and prone to human error. Linux open-source HA extensions require a high degree of technical skill, creating complexity and reliability issues that challenge most operators. "Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	SIOS Clustering for Linux SIOS is a high availability company that has spent the past 20 years focused on delivering HA that is specifically designed for SAP, SQL, Linux, Oracle, and other applications. Its experience is built into its product, and installation and configuration take a fraction of the time and cost when compared to custom scripting with the Linux distributions. In addition, SIOS tests and validates new versions of operating systems and applications so its customers don't have to. When a customer calls SIOS for support, they are connected to a high availability expert – someone who only focuses on HA and has been doing so for a very long time. In Linux the most used software is **Pacemaker** Pacemaker overview The High Availability Add-On cluster infrastructure provides the basic functions for a group of computers (called nodes or members) to work together as a cluster. Once a cluster is formed using the cluster infrastructure, you can use other components to suit your clustering needs (for example, setting up a cluster for sharing files on a GFS2 file system or setting up service failover). The cluster infrastructure performs the following functions: • Cluster management. • Lock management. • Fencing. • Cluster configuration management. "Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	Overview *High availability* is the ability of an IT system to be accessible and reliable nearly 100% of the time, eliminating or minimizing downtime. It combines two concepts to determine if an IT system is meeting its operational performance level: that a given service or server is accessible–or available–almost 100% of the time without downtime, and that the service or server performs to reasonable expectations for an established time period. High availability is more than hitting an uptime service level agreement (SLA), or the expectations set between a service provider and client. It is about truly resilient, reliable, and well-functioning systems. "Red Hat – What is high availability."[14]

[14] Available at https://www.redhat.com/en/topics/linux/what-is-high-availability.

Claim 1	Exemplary Dell Servers
	What are high-availability clusters? High-availability architectures run active failover clusters, so there is built-in redundancy and failover and—hopefully—zero downtime. Within the cluster, nodes are monitored not just for availability, but for overall performance of applications, services, and network. Because there is shared storage, there is no data loss if a node goes down, because all cluster nodes work from the same data source. Load balancing can be used to manage traffic for best performance. Outside those broad characteristics, high-availability clusters can be designed for more specialized activities, depending on the priorities and activities within the IT infrastructure. The Red Hat Enterprise Linux High Availability Add-on, for example, has four default configurations: • **High availability**: focuses on uptime and availability • **High performance**: for high speed, concurrent operations • **Load balancing**: for cost-effective scalability • **Storage**: for resilient data management In real-life environments, the high-availability systems would incorporate aspects of those focus elements. "Red Hat – What is high availability."

Claim 1	Exemplary Dell Servers

High Availability Add-On

Red Hat® Enterprise Linux® provides the ability to create managed, highly available clusters with groups of RHEL servers. This Add-On can be configured to manage most applications (both off-the-shelf and custom) and provides a wide range of configuration options to fit most requirements. This Add-On includes:

- The Pacemaker cluster resource manager supports up to 32 (x86) nodes.
- Corosync and Kronosnet for managing network communication between cluster nodes.
- Booth and QDevice/QNetd for managing multisite and stretch clusters.
- A Pacemaker command line interface (pcs) and GUI (pcsd).
- A collection of resource agents for many commonly used applications as well as support for user-created resource agents, management of systemd services, and containers.
- A collection of fencing agents for use on bare metal servers, virtual machines, and cloud platforms.
- For IBM POWER and/or Z, please contact your sales representative.

Red Hat® Enterprise Linux® targets general-purpose clustering workloads. Its capabilities include support for a wide variety of application configurations, including active/passive, active/active & primary/secondary. The clustering software maintains high availability by automatically restarting and/or moving failing services (and their dependencies) to other functioning nodes.

Some well-known applications that run well with RHEL High Availability include SAP, Apache, PostgreSQL, and DB2. In addition, most custom applications can be managed by RHEL High Availability by using their systemd startup script or by writing a simple custom resource agent.

"Red Hat – High Availability Add-On."[15]

[15] Available at https://www.redhat.com/en/store/high-availability-add.

Claim 1	Exemplary Dell Servers
[b] a second computer that is able to perform at least one of scale-up and scale-down when operating as a standby system of the redundant system; and	The Exemplary Dell Servers comprise a second computer that is able to perform at least one of scale-up and scale-down when operating as a standby system of the redundant system. For example, any other computer in a cluster could be the "second computer" that is able to perform at least one of scale-up and scale-down when operating as a standby system of the redundant system. Many options are available, including scaling down a host within a cluster that loses its management network connection (referred to as "host isolation" response), and scaling up a new (e.g., non-isolated) host. *See, e.g.,*: Host Isolation Response Host isolation response determines what happens when a host in a vSphere HA cluster loses its management network connections, but continues to run. You can use the isolation response to have vSphere HA power off virtual machines that are running on an isolated host and restart them on a non-isolated host. Host isolation responses require that Host Monitoring Status is activated. If Host Monitoring Status is deactivated, host isolation responses are also suspended. A host determines that it is isolated when it is unable to communicate with the agents running on the other hosts, and it is unable to ping its isolation addresses. The host then executes its isolation response. The responses are Power off and restart VMs or Shutdown and restart VMs. You can customize this property for individual virtual machines. "Determining Responses to Host Issues."[16] *See also* "Host Failure Types," for other host failure scenarios (e.g., "failure," "isolation," and "partition") that also result in a scale-up or scale-down operation.

[16] Available at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-FA8B166D-A5F5-47D3-840E-68996507A95B.html.

Claim 1	Exemplary Dell Servers

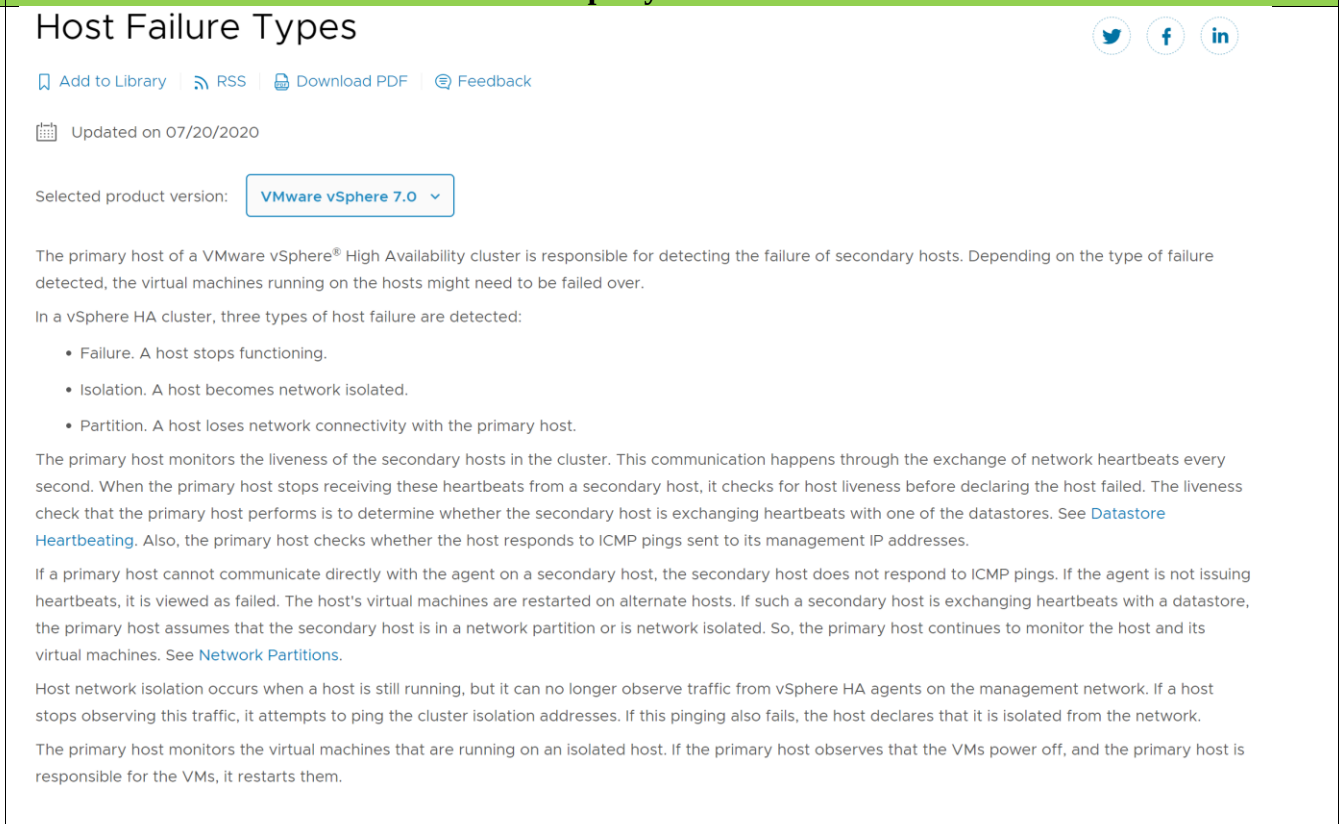

Host Failure Types

Add to Library | RSS | Download PDF | Feedback

Updated on 07/20/2020

Selected product version: VMware vSphere 7.0

The primary host of a VMware vSphere® High Availability cluster is responsible for detecting the failure of secondary hosts. Depending on the type of failure detected, the virtual machines running on the hosts might need to be failed over.

In a vSphere HA cluster, three types of host failure are detected:

- Failure. A host stops functioning.
- Isolation. A host becomes network isolated.
- Partition. A host loses network connectivity with the primary host.

The primary host monitors the liveness of the secondary hosts in the cluster. This communication happens through the exchange of network heartbeats every second. When the primary host stops receiving these heartbeats from a secondary host, it checks for host liveness before declaring the host failed. The liveness check that the primary host performs is to determine whether the secondary host is exchanging heartbeats with one of the datastores. See Datastore Heartbeating. Also, the primary host checks whether the host responds to ICMP pings sent to its management IP addresses.

If a primary host cannot communicate directly with the agent on a secondary host, the secondary host does not respond to ICMP pings. If the agent is not issuing heartbeats, it is viewed as failed. The host's virtual machines are restarted on alternate hosts. If such a secondary host is exchanging heartbeats with a datastore, the primary host assumes that the secondary host is in a network partition or is network isolated. So, the primary host continues to monitor the host and its virtual machines. See Network Partitions.

Host network isolation occurs when a host is still running, but it can no longer observe traffic from vSphere HA agents on the management network. If a host stops observing this traffic, it attempts to ping the cluster isolation addresses. If this pinging also fails, the host declares that it is isolated from the network.

The primary host monitors the virtual machines that are running on an isolated host. If the primary host observes that the VMs power off, and the primary host is responsible for the VMs, it restarts them.

"Host Failure Types."[17]

[17] Available at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-4ED552C3-0291-4553-A46A-290EF883BC8D.html.

Claim 1	Exemplary Dell Servers
	Server (3) FaultType: Microsoft.Health.FaultType.Server.Down • Severity: Critical • Reason: "*The server cannot be reached.*" • RecommendedAction: "*Start or replace server.*" FaultType: Microsoft.Health.FaultType.Server.Isolated • Severity: Critical • Reason: "*The server is isolated from the cluster due to connectivity issues.*" • RecommendedAction: "*If isolation persists, check the network(s) or migrate workloads to other nodes.*" FaultType: Microsoft.Health.FaultType.Server.Quarantined • Severity: Critical • Reason: "*The server is quarantined by the cluster due to recurring failures.*" • RecommendedAction: "*Replace the server or fix the network.*" "Health Service Faults."[18]

[18] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/health-service-faults.

Claim 1	Exemplary Dell Servers

	Virtual Machine Resiliency
	
	Compute Resiliency Windows Server 2016 includes increased virtual machines compute resiliency to help reduce intra-cluster communication issues in your compute cluster as follows:
	
	• **Resiliency options available for virtual machines:** You can now configure virtual machine resiliency options that define behavior of the virtual machines during transient failures:
	
	  ○ **Resiliency Level:** Helps you define how the transient failures are handled.
	
	  ○ **Resiliency Period:** Helps you define how long all the virtual machines are allowed to run isolated.
	
	• **Quarantine of unhealthy nodes:** Unhealthy nodes are quarantined and are no longer allowed to join the cluster. This prevents flapping nodes from negatively effecting other nodes and the overall cluster.
	
	"What's New in Failover Clustering."

Claim 1	Exemplary Dell Servers

High Availability Clusters

High-availability clusters (also known as **HA clusters** , **failover clusters** or **Metroclusters Active/Active**) are groups of computers that support server applications that can be reliably utilised with a minimum amount of down-time.

They operate by using high availability software to harness redundant computers in groups or clusters that provide continued service when system components fail.

Without clustering, if a server running a particular application crashes, the application will be unavailable until the crashed server is fixed. HA clustering remedies this situation by detecting hardware/software faults, and immediately restarting the application on another system without requiring administrative intervention, a process known as failover.

As part of this process, clustering software may configure the node before starting the application on it. For example, appropriate file systems may need to be imported and mounted, network hardware may have to be configured, and some supporting applications may need to be running as well.

HA clusters are often used for critical databases, file sharing on a network, business applications, and customer services such as electronic commerce websites.

"Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers
	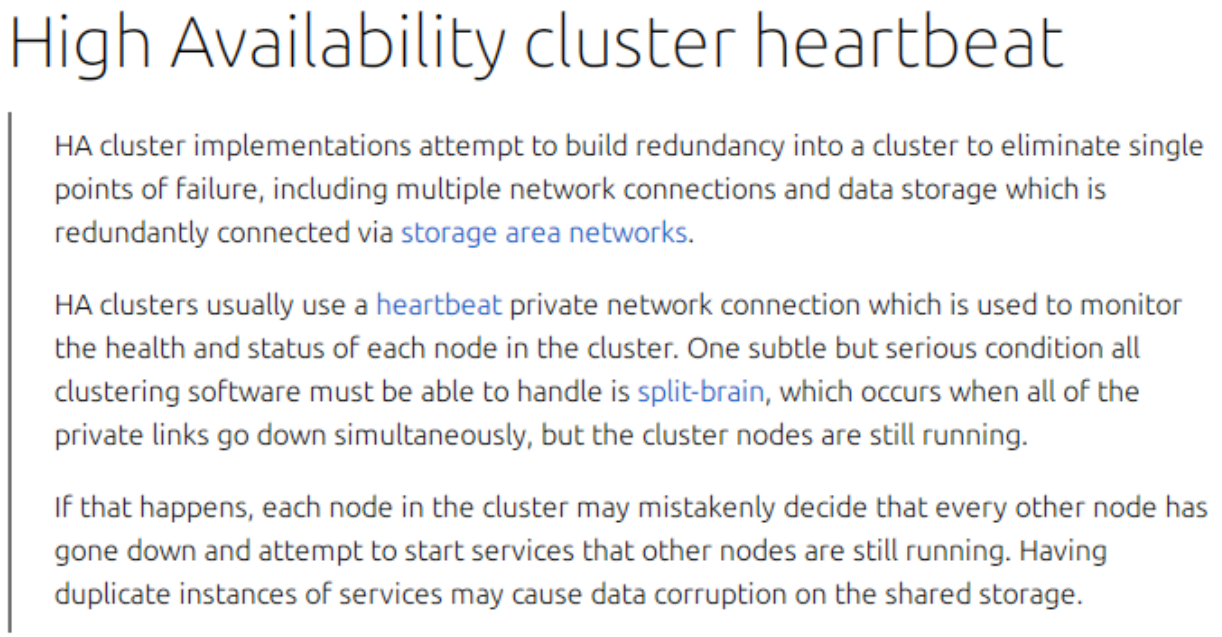 "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Linux Clustering Concepts

A cluster is a group of computers (nodes) which work together to provide a shared solution. At a high level, a cluster can be viewed as having three parts (often defined as cluster stack).

Basic concepts

- **Resources:** These are the reason for the cluster's being the services that need to be kept highly available.
- **Resource Agents:** These are scripts operating system components that start, stop, and monitor resources, given a set of resource parameters.
- **Fence Agents:** These are scripts that execute a node fencing actions, given a target and fence device.
- **Fencing:** The ability to disable nodes.
- **Quorum:** Encapsulates the ability to determine whether the cluster can continue to operate safely or not.

Cluster types

The four types are as follows:

- **High Availability (HA):** Used for Fault Tolerance to keep server services available to employees or customers.
- **Load Balancing:** Balances the load between multiple systems when a service needs to be available to numerous systems at once (can be used for other three types of Clusters).
- **Distributed:** Jobs will be managed by different systems.
- **Parallel (Beowulf):** Jobs are managed by multiple processors on multiple systems.

As well by the configuration there are several types of clusters:

- **Manual clustering:** Lets you classify, merge, split clusters manually if the output of the automatic spike sorting algorithms are not satisfactory.
- **Merging clusters:** When multiple clusters seem to correspond to the same unit.
- **Splitting clusters:** You can create a new cluster by drawing a polygon around a set of spikes in the feature view, the amplitude view, the template amplitude view, or the spike attribute views.

"Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	Reasons to have a cluster There are different reasons for using a cluster. We use them to provide a resilient, highly available backend for solutions. This is primarily for backend services like MySQL, NFS or Redis where the service can be behind a front end of web servers. High availability is provided via automatic failover - if there is a failure on a cluster node, the clustered services running on that node will automatically relocate away to a node that is running properly. Clusters will generally use shared storage (SAN) so that data will be persistent when services move between nodes. In the event if a node becomes unresponsive, typically that node will be rebooted (fenced) by other nodes in order to preserve data integrity of the shared storage and floating IP ownership. **Linux Open Source High Availability Clustering** Some Linux operating system vendors offer clustering software, such as SUSE Linux HAE; Red Hat Enterprise Linux (RHEL); and Oracle Real Application Clusters (RAC). While they allow you to create a failover cluster, they present a variety of challenges, thist is highly manual and prone to human error. Linux open-source HA extensions require a high degree of technical skill, creating complexity and reliability issues that challenge most operators. "Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	What are high-availability clusters? High-availability architectures run active failover clusters, so there is built-in redundancy and failover and—hopefully—zero downtime. Within the cluster, nodes are monitored not just for availability, but for overall performance of applications, services, and network. Because there is shared storage, there is no data loss if a node goes down, because all cluster nodes work from the same data source. Load balancing can be used to manage traffic for best performance. Outside those broad characteristics, high-availability clusters can be designed for more specialized activities, depending on the priorities and activities within the IT infrastructure. The Red Hat Enterprise Linux High Availability Add-on, for example, has four default configurations: • **High availability**: focuses on uptime and availability • **High performance**: for high speed, concurrent operations • **Load balancing**: for cost-effective scalability • **Storage**: for resilient data management In real-life environments, the high-availability systems would incorporate aspects of those focus elements. "Red Hat – What is high availability."
[c] a controller that issues an instruction to the second computer operating as the standby system to perform	The Exemplary Dell Servers comprise a controller that issues an instruction to the second computer operating as the standby system to perform the scale-up or the scale-down, when the active system needs to be scaled-up or scaled-down.

Claim 1	Exemplary Dell Servers
the scale-up or the scale-down, when the active system needs to be scaled-up or scaled-down,	For example, the vCenter Server, or the primary host acting as the vCenter Server management interface, includes a controller that issues an instruction to the second computer operating as the standby system to perform the scale-up or the scale-down, when the active system needs to be scaled-up or scaled-down. For example, a "cluster health state" is maintained and used to determine which nodes to perform the scale-up or scale-down. *See, e.g.,*: 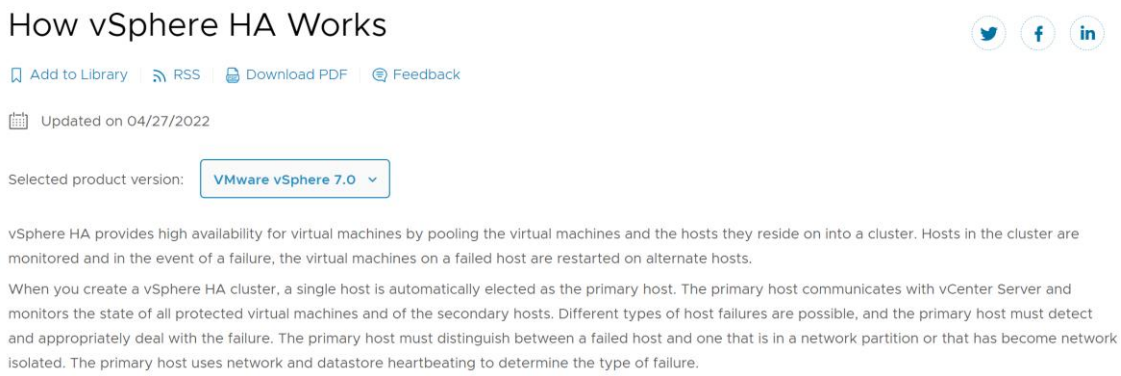 "How vSphere HA Works."

Claim 1	Exemplary Dell Servers

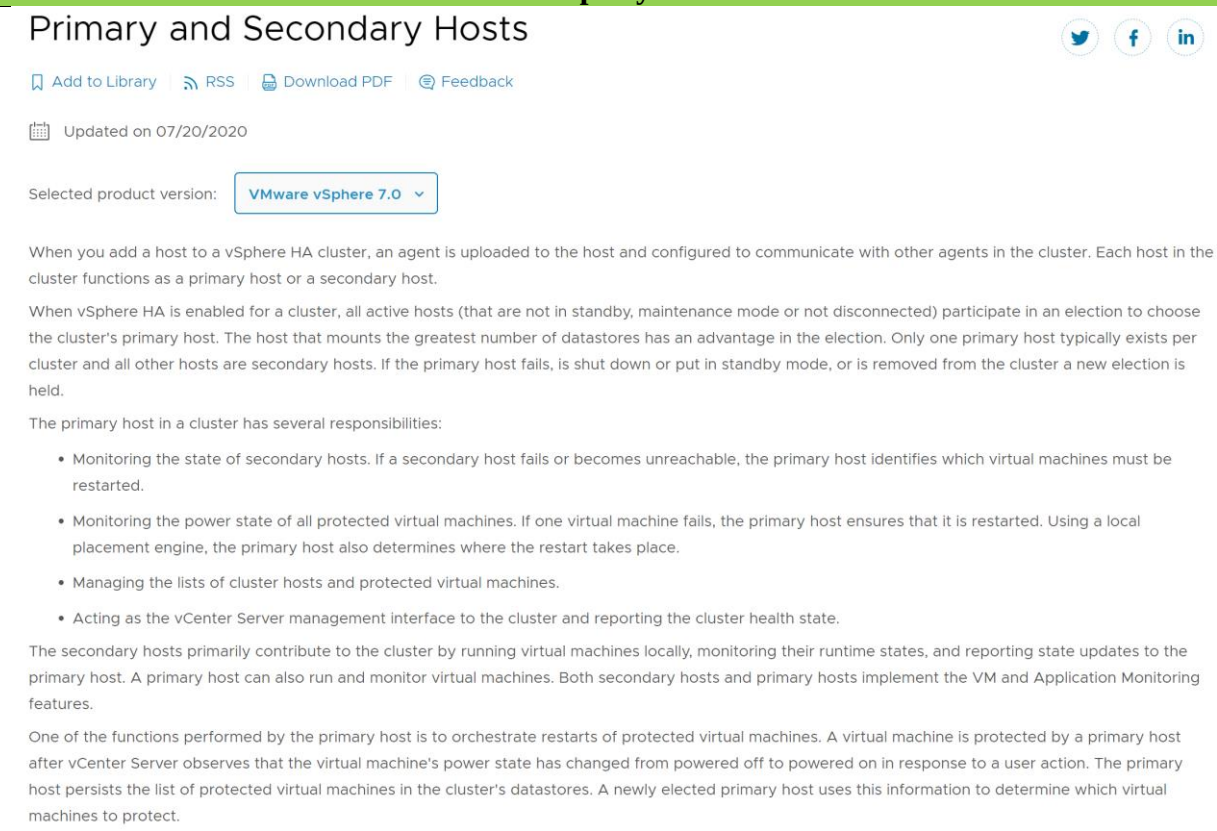

Primary and Secondary Hosts

Add to Library | RSS | Download PDF | Feedback

Updated on 07/20/2020

Selected product version: VMware vSphere 7.0

When you add a host to a vSphere HA cluster, an agent is uploaded to the host and configured to communicate with other agents in the cluster. Each host in the cluster functions as a primary host or a secondary host.

When vSphere HA is enabled for a cluster, all active hosts (that are not in standby, maintenance mode or not disconnected) participate in an election to choose the cluster's primary host. The host that mounts the greatest number of datastores has an advantage in the election. Only one primary host typically exists per cluster and all other hosts are secondary hosts. If the primary host fails, is shut down or put in standby mode, or is removed from the cluster a new election is held.

The primary host in a cluster has several responsibilities:

- Monitoring the state of secondary hosts. If a secondary host fails or becomes unreachable, the primary host identifies which virtual machines must be restarted.
- Monitoring the power state of all protected virtual machines. If one virtual machine fails, the primary host ensures that it is restarted. Using a local placement engine, the primary host also determines where the restart takes place.
- Managing the lists of cluster hosts and protected virtual machines.
- Acting as the vCenter Server management interface to the cluster and reporting the cluster health state.

The secondary hosts primarily contribute to the cluster by running virtual machines locally, monitoring their runtime states, and reporting state updates to the primary host. A primary host can also run and monitor virtual machines. Both secondary hosts and primary hosts implement the VM and Application Monitoring features.

One of the functions performed by the primary host is to orchestrate restarts of protected virtual machines. A virtual machine is protected by a primary host after vCenter Server observes that the virtual machine's power state has changed from powered off to powered on in response to a user action. The primary host persists the list of protected virtual machines in the cluster's datastores. A newly elected primary host uses this information to determine which virtual machines to protect.

"Primary and Secondary Hosts."[19]

For example, the determination whether to perform a scale-up or scale-down is implemented in part using "heartbeats" that monitor for host failures, including failed, isolated, and partitioned hosts.

[19] Available at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-B9ACCE9B-A14D-4B2F-87EF-4B0A67ABDF58.html.

Claim 1	Exemplary Dell Servers
	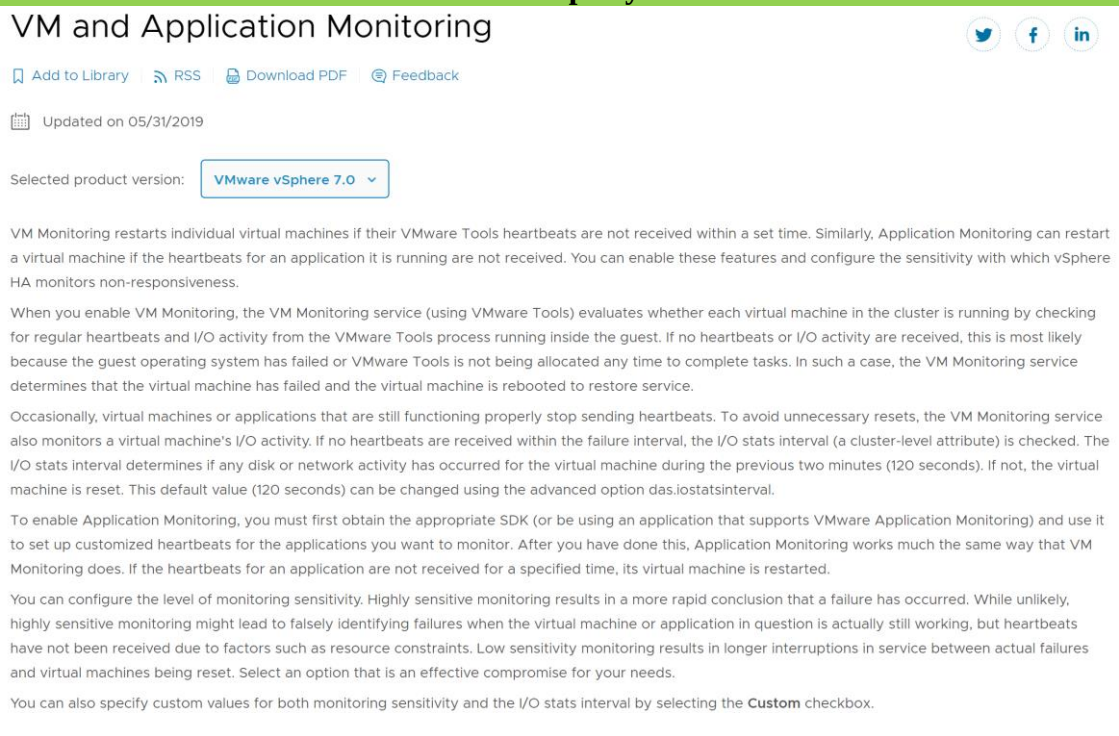 "VM and Application Monitoring."[20]

[20] Available at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-62B80D7A-C764-40CB-AE59-752DA6AD78E7.html.

Claim 1	Exemplary Dell Servers
	> Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2 A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service. Failover Clustering has many practical applications, including: • Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines • Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V "Failover Clustering Overview."

Claim 1	Exemplary Dell Servers
	# Health Service faults Article • 08/25/2021 > Applies to: Windows Server 2016 ## What are faults The Health Service constantly monitors your Storage Spaces Direct cluster to detect problems and generate "faults". One new cmdlet displays any current faults, allowing you to easily verify the health of your deployment without looking at every entity or feature in turn. Faults are designed to be precise, easy to understand, and actionable. Each fault contains five important fields: • Severity • Description of the problem • Recommended next step(s) to address the problem • Identifying information for the faulting entity • Its physical location (if applicable) For example, here is a typical fault: ``` Severity: MINOR Reason: Connectivity has been lost to the physical disk. Recommendation: Check that the physical disk is working and properly connected. Part: Manufacturer Contoso, Model XYZ9000, Serial 123456789 Location: Seattle DC, Rack B07, Node 4, Slot 11 ``` "Health Service Faults."

Claim 1	Exemplary Dell Servers

High Availability Clusters

> **High-availability clusters** (also known as **HA clusters** , **failover clusters** or **Metroclusters Active/Active**) are groups of computers that support server applications that can be reliably utilised with a minimum amount of down-time.
>
> They operate by using high availability software to harness redundant computers in groups or clusters that provide continued service when system components fail.
>
> Without clustering, if a server running a particular application crashes, the application will be unavailable until the crashed server is fixed. HA clustering remedies this situation by detecting hardware/software faults, and immediately restarting the application on another system without requiring administrative intervention, a process known as failover.
>
> As part of this process, clustering software may configure the node before starting the application on it. For example, appropriate file systems may need to be imported and mounted, network hardware may have to be configured, and some supporting applications may need to be running as well.
>
> HA clusters are often used for critical databases, file sharing on a network, business applications, and customer services such as electronic commerce websites.

"Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers
	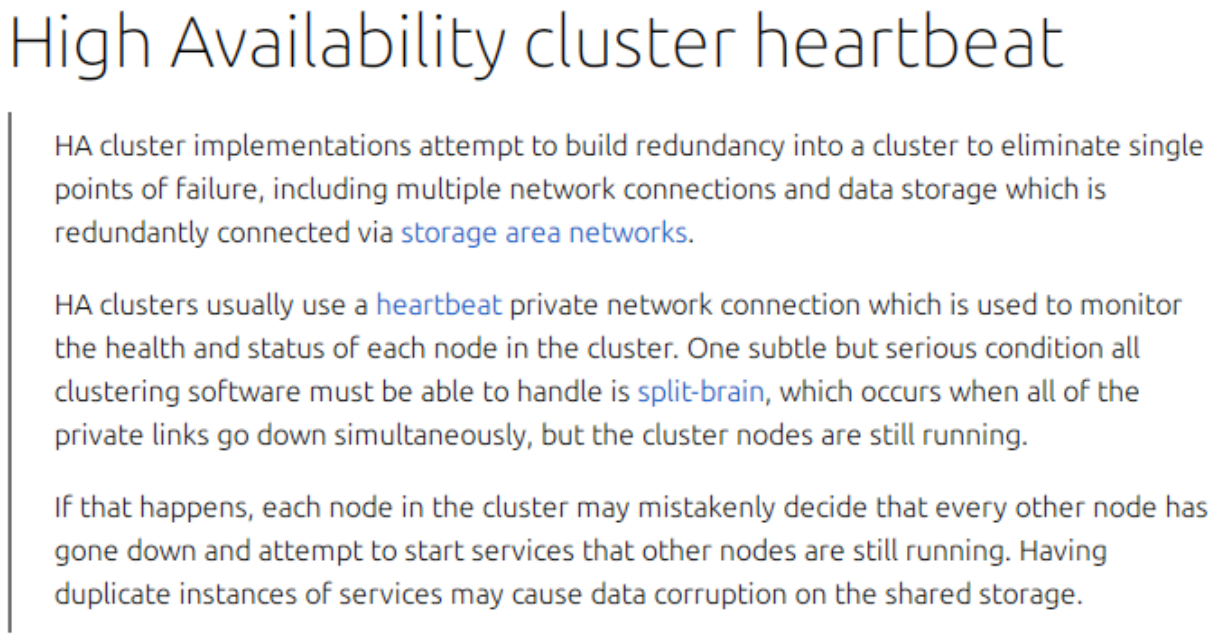 "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Linux Clustering Concepts

A cluster is a group of computers (nodes) which work together to provide a shared solution. At a high level, a cluster can be viewed as having three parts (often defined as cluster stack).

Basic concepts

- **Resources:** These are the reason for the cluster's being the services that need to be kept highly available.
- **Resource Agents:** These are scripts operating system components that start, stop, and monitor resources, given a set of resource parameters.
- **Fence Agents:** These are scripts that execute a node fencing actions, given a target and fence device.
- **Fencing:** The ability to disable nodes.
- **Quorum:** Encapsulates the ability to determine whether the cluster can continue to operate safely or not.

Cluster types

The four types are as follows:

- **High Availability (HA):** Used for Fault Tolerance to keep server services available to employees or customers.
- **Load Balancing:** Balances the load between multiple systems when a service needs to be available to numerous systems at once (can be used for other three types of Clusters).
- **Distributed:** Jobs will be managed by different systems.
- **Parallel (Beowulf):** Jobs are managed by multiple processors on multiple systems.

As well by the configuration there are several types of clusters:

- **Manual clustering:** Lets you classify, merge, split clusters manually if the output of the automatic spike sorting algorithms are not satisfactory.
- **Merging clusters:** When multiple clusters seem to correspond to the same unit.
- **Splitting clusters:** You can create a new cluster by drawing a polygon around a set of spikes in the feature view, the amplitude view, the template amplitude view, or the spike attribute views.

"Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	Reasons to have a cluster There are different reasons for using a cluster. We use them to provide a resilient, highly available backend for solutions. This is primarily for backend services like MySQL, NFS or Redis where the service can be behind a front end of web servers. High availability is provided via automatic failover - if there is a failure on a cluster node, the clustered services running on that node will automatically relocate away to a node that is running properly. Clusters will generally use shared storage (SAN) so that data will be persistent when services move between nodes. In the event if a node becomes unresponsive, typically that node will be rebooted (fenced) by other nodes in order to preserve data integrity of the shared storage and floating IP ownership. **Linux Open Source High Availability Clustering** Some Linux operating system vendors offer clustering software, such as SUSE Linux HAE; Red Hat Enterprise Linux (RHEL); and Oracle Real Application Clusters (RAC). While they allow you to create a failover cluster, they present a variety of challenges, thist is highly manual and prone to human error. Linux open-source HA extensions require a high degree of technical skill, creating complexity and reliability issues that challenge most operators. "Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	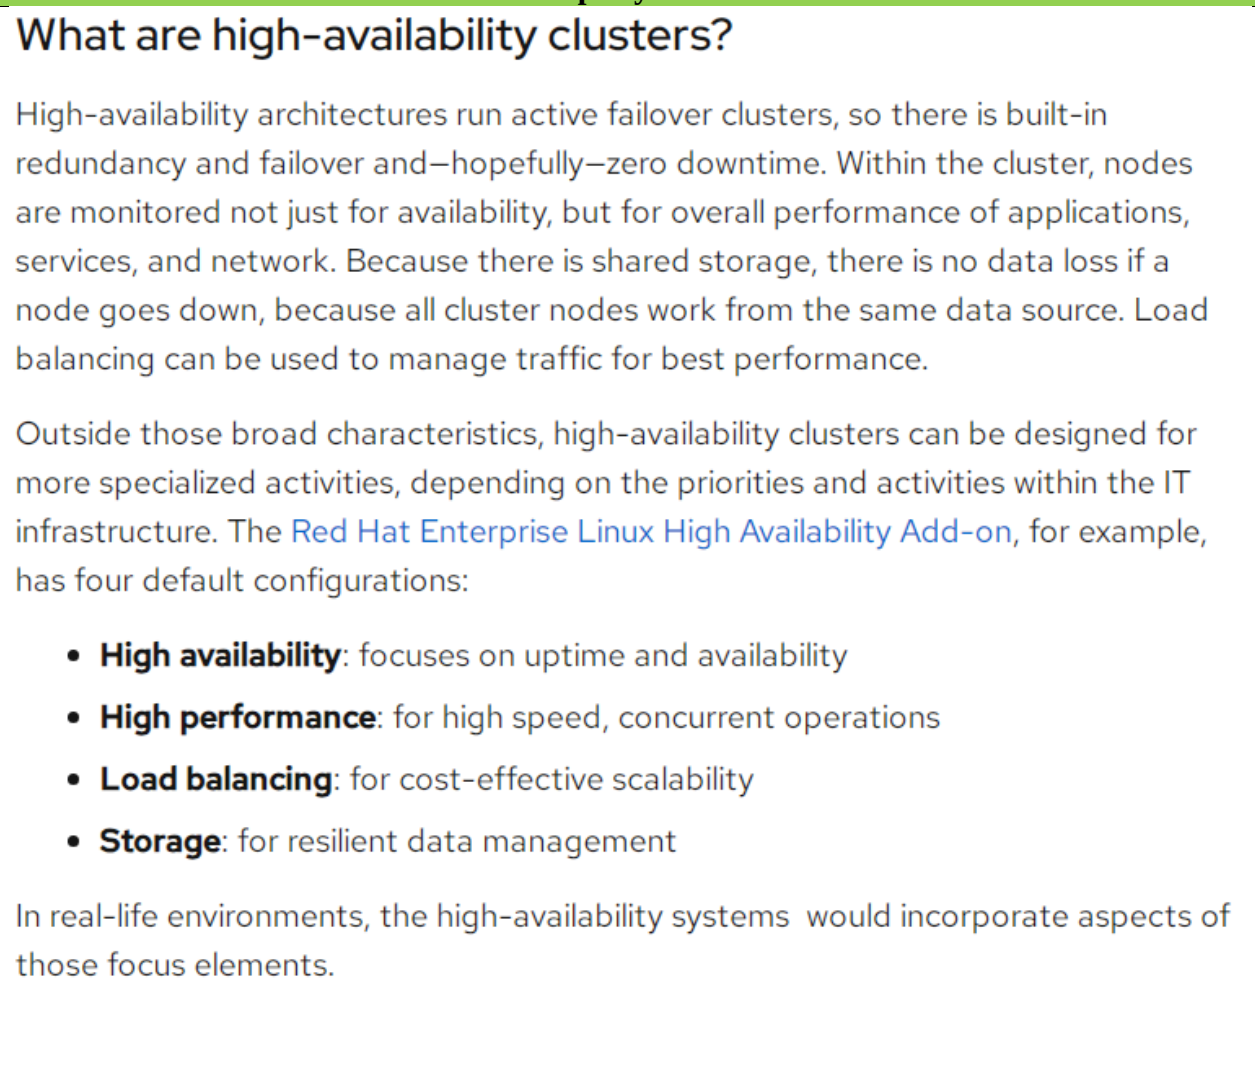 "Red Hat – What is high availability."
[d] wherein the second computer operating as the standby system, responsive	The second computer in the Exemplary Dell Servers, operating as the standby system, responsive to the instruction, in case of performing the scale-up, increases the number of virtual CPUs (Central Processing Units) included in the second computer and allocates one or more processes to one or

Claim 1	Exemplary Dell Servers
to the instruction, in case of performing the scale-up, increases the number of virtual CPUs (Central Processing Units) included in the second computer and allocates one or more processes to one or more virtual CPUs added, while in case of performing the scale-down, the second computer decreases the number of virtual CPUs included in the second computer and releases allocation of one or more processes allocated to one or more virtual CPUs deleted, and transmits a completion notification to the controller when the scale-up or the scale-down is completed, and	more virtual CPUs added, while in case of performing the scale-down, the second computer decreases the number of virtual CPUs included in the second computer and releases allocation of one or more processes allocated to one or more virtual CPUs deleted, and transmits a completion notification to the controller when the scale-up or the scale-down is completed. For example, VMs may be terminated and restarted (optionally in priority order) on the second computer. *See, e.g.,*: Determining Responses to Host Issues Add to Library | RSS | Download PDF | Feedback Updated on 08/23/2022 Selected product version: VMware vSphere 7.0 If a host fails and its virtual machines must be restarted, you can control the order in which the virtual machines are restarted with the VM restart priority setting. You can also configure how vSphere HA responds if hosts lose management network connectivity with other hosts by using the host isolation response setting. Other factors are also considered when vSphere HA restarts a virtual machine after a failure. The following settings apply to all virtual machines in the cluster in the case of a host failure or isolation. You can also configure exceptions for specific virtual machines. See Customize an Individual Virtual Machine. "Determining Responses to Host Issues." For example, a variety of factors are considered when performing the scale-up or scale-down, including a maximum number of allowed VMs or the number of in-use VCPUs.

Claim 1	Exemplary Dell Servers
	Factors Considered for Virtual Machine Restarts After a failure, the cluster's primary host attempts to restart affected virtual machines by identifying a host that can power them on. When choosing such a host, the primary host considers a number of factors. **File accessibility** Before a virtual machine can be started, its files must be accessible from one of the active cluster hosts that the primary can communicate with over the network **Virtual machine and host compatibility** If there are accessible hosts, the virtual machine must be compatible with at least one of them. The compatibility set for a virtual machine includes the effect of any required VM-Host affinity rules. For example, if a rule only permits a virtual machine to run on two hosts, it is considered for placement on those two hosts. **Resource reservations** Of the hosts that the virtual machine can run on, at least one must have sufficient unreserved capacity to meet the memory overhead of the virtual machine and any resource reservations. Four types of reservations are considered: CPU, Memory, vNIC, and Virtual flash. Also, sufficient network ports must be available to power on the virtual machine. **Host limits** In addition to resource reservations, a virtual machine can only be placed on a host if doing so does not violate the maximum number of allowed virtual machines or the number of in-use vCPUs. "Determining Responses to Host Issues."

Claim 1	Exemplary Dell Servers
	> Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2 A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service. Failover Clustering has many practical applications, including: • Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines • Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V "Failover Clustering Overview."

Claim 1	Exemplary Dell Servers
	# What's new in Windows Server 2019 and Azure Stack HCI • **Cluster sets** (Applies only to Windows Server 2019) Cluster sets enable you to increase the number of servers in a single software-defined datacenter (SDDC) solution beyond the current limits of a cluster. This is accomplished by grouping multiple clusters into a cluster set, a loosely coupled grouping of multiple failover clusters: compute, storage and hyper-converged. With cluster sets, you can move online virtual machines (live migrate) between clusters within the cluster set. For more info, see Cluster sets. • **Azure-aware clusters** Failover clusters now automatically detect when they're running in Azure IaaS virtual machines and optimize the configuration to provide proactive failover and logging of Azure planned maintenance events to achieve the highest levels of availability. Deployment is also simplified by removing the need to configure the load balancer with Distributed Network Name for cluster name. "What's New in Failover Clustering."

Claim 1	Exemplary Dell Servers

	Virtual Machine Resiliency
	
	Compute Resiliency Windows Server 2016 includes increased virtual machines compute resiliency to help reduce intra-cluster communication issues in your compute cluster as follows:
	
	• **Resiliency options available for virtual machines:** You can now configure virtual machine resiliency options that define behavior of the virtual machines during transient failures:
	
	  ○ **Resiliency Level:** Helps you define how the transient failures are handled.
	
	  ○ **Resiliency Period:** Helps you define how long all the virtual machines are allowed to run isolated.
	
	• **Quarantine of unhealthy nodes:** Unhealthy nodes are quarantined and are no longer allowed to join the cluster. This prevents flapping nodes from negatively effecting other nodes and the overall cluster.
	
	"What's New in Failover Clustering."

Claim 1	Exemplary Dell Servers
	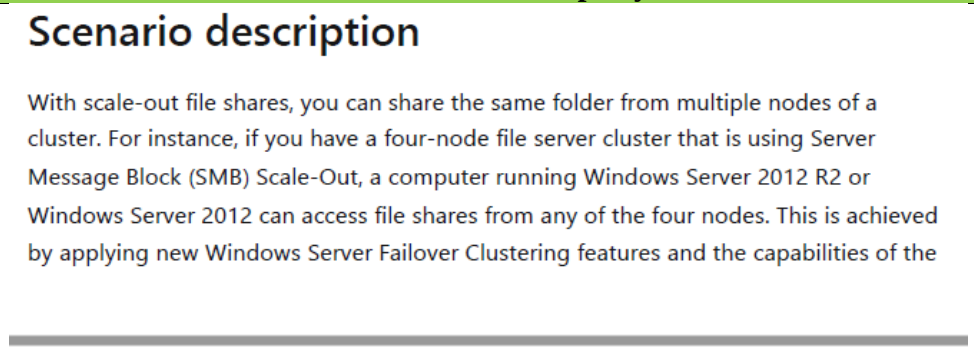 "Scale-Out File Server for application data overview."[21]

[21] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/sofs-overview.

Claim 1	Exemplary Dell Servers
	Benefits Cluster sets offer the following benefits: • Significantly increases the supported SDDC cloud scale for running highly available virtual machines (VMs) by combining multiple smaller clusters into a single large fabric, while keeping the software fault boundary to a single cluster. You can easily migrate VMs across the cluster set. • Increased resiliency. Having four 4-node clusters in a cluster set gives you better resiliency than a single 16-node cluster in that multiple compute nodes can go down and production remains intact. • Management of failover cluster lifecycle, including onboarding and retiring clusters, without impacting tenant VM availability. • VM flexibility across individual clusters and a present a unified storage namespace. • Easily change the compute-to-storage workload ratio in your hyper-converged environment. • Benefit from Azure-like Fault Domains and Availability sets across individual clusters in initial VM placement and subsequent migration. • Can use even if compute and storage hardware between cluster nodes isn't identical. • Live migration of VMs between clusters. • Azure-like availability sets and fault domains across multiple clusters.

Claim 1	Exemplary Dell Servers

	"Deploy a cluster set."[22]
	
	When complete, you are shown which cluster node the VM was deployed on. For the above example, it would show as:
	
	`State : Running`
	`ComputerName : 1-S2D2`
	
	"Deploy a cluster set."

[22] Available at https://learn.microsoft.com/en-us/windows-server/failover-clustering/cluster-set.

Claim 1	Exemplary Dell Servers

	Once the VM is created, it is displayed in Hyper-V manager on the specific node specified. To add it as a cluster set VM and add it to the cluster, use this command:
	
	PowerShell
	
	`Register-ClusterSetVM -CimSession CSMASTER -MemberName $targetnode.Member -VMName CSVM1`
	
	When complete, the output is:
	
	`Id VMName State MemberName PSComputerName`
	`-- ------ ----- ---------- --------------`
	`1 CSVM1 On CLUSTER1 CSMASTER`
	
	"Deploy a cluster set."

Claim 1	Exemplary Dell Servers
	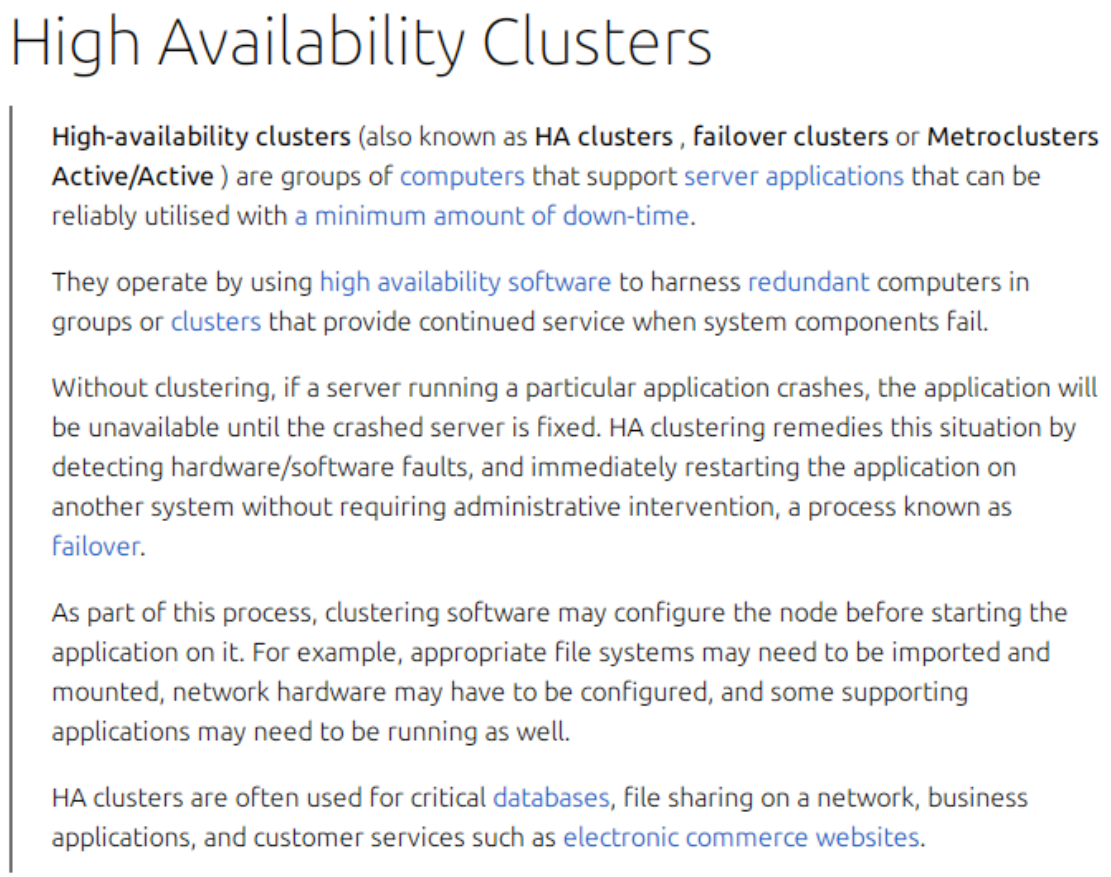 "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers
	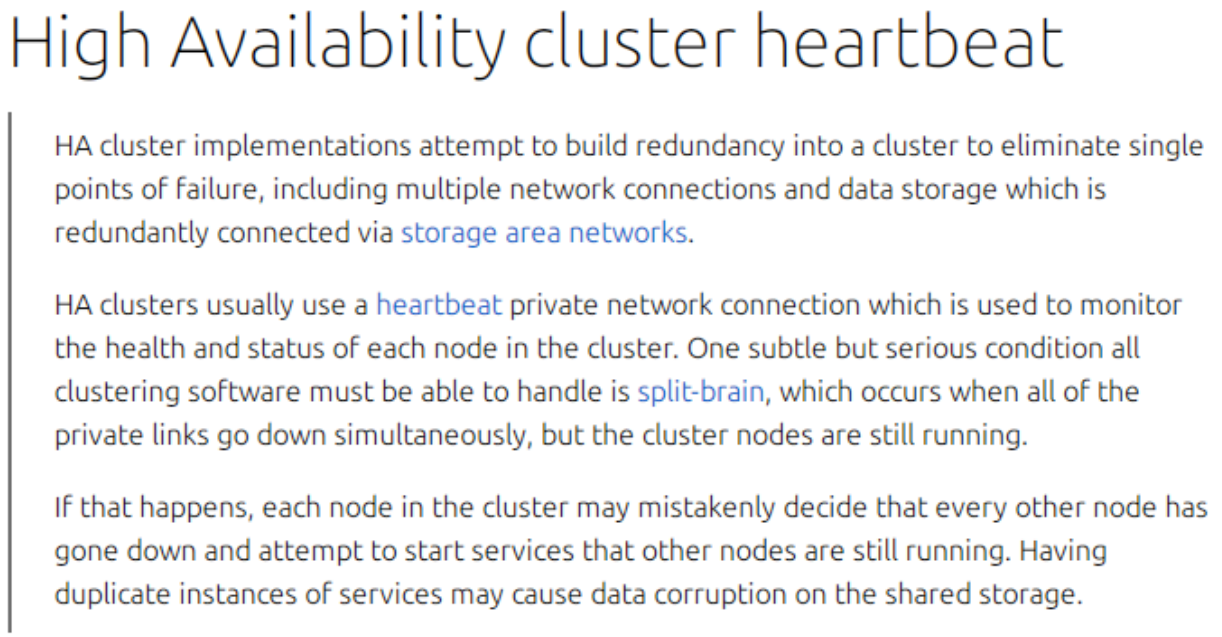 "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Linux Clustering Concepts

A cluster is a group of computers (nodes) which work together to provide a shared solution. At a high level, a cluster can be viewed as having three parts (often defined as cluster stack).

Basic concepts

- **Resources:** These are the reason for the cluster's being the services that need to be kept highly available.
- **Resource Agents:** These are scripts operating system components that start, stop, and monitor resources, given a set of resource parameters.
- **Fence Agents:** These are scripts that execute a node fencing actions, given a target and fence device.
- **Fencing:** The ability to disable nodes.
- **Quorum:** Encapsulates the ability to determine whether the cluster can continue to operate safely or not.

Cluster types

The four types are as follows:

- **High Availability (HA):** Used for Fault Tolerance to keep server services available to employees or customers.
- **Load Balancing:** Balances the load between multiple systems when a service needs to be available to numerous systems at once (can be used for other three types of Clusters).
- **Distributed:** Jobs will be managed by different systems.
- **Parallel (Beowulf):** Jobs are managed by multiple processors on multiple systems.

As well by the configuration there are several types of clusters:

- **Manual clustering:** Lets you classify, merge, split clusters manually if the output of the automatic spike sorting algorithms are not satisfactory.
- **Merging clusters:** When multiple clusters seem to correspond to the same unit.
- **Splitting clusters:** You can create a new cluster by drawing a polygon around a set of spikes in the feature view, the amplitude view, the template amplitude view, or the spike attribute views.

"Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	Reasons to have a cluster There are different reasons for using a cluster. We use them to provide a resilient, highly available backend for solutions. This is primarily for backend services like MySQL, NFS or Redis where the service can be behind a front end of web servers. High availability is provided via automatic failover - if there is a failure on a cluster node, the clustered services running on that node will automatically relocate away to a node that is running properly. Clusters will generally use shared storage (SAN) so that data will be persistent when services move between nodes. In the event if a node becomes unresponsive, typically that node will be rebooted (fenced) by other nodes in order to preserve data integrity of the shared storage and floating IP ownership. **Linux Open Source High Availability Clustering** Some Linux operating system vendors offer clustering software, such as SUSE Linux HAE; Red Hat Enterprise Linux (RHEL); and Oracle Real Application Clusters (RAC). While they allow you to create a failover cluster, they present a variety of challenges, thist is highly manual and prone to human error. Linux open-source HA extensions require a high degree of technical skill, creating complexity and reliability issues that challenge most operators. "Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	What are high-availability clusters? High-availability architectures run active failover clusters, so there is built-in redundancy and failover and—hopefully—zero downtime. Within the cluster, nodes are monitored not just for availability, but for overall performance of applications, services, and network. Because there is shared storage, there is no data loss if a node goes down, because all cluster nodes work from the same data source. Load balancing can be used to manage traffic for best performance. Outside those broad characteristics, high-availability clusters can be designed for more specialized activities, depending on the priorities and activities within the IT infrastructure. The Red Hat Enterprise Linux High Availability Add-on, for example, has four default configurations: • **High availability**: focuses on uptime and availability • **High performance**: for high speed, concurrent operations • **Load balancing**: for cost-effective scalability • **Storage**: for resilient data management In real-life environments, the high-availability systems would incorporate aspects of those focus elements. "Red Hat – What is high availability."
[e] wherein, upon reception of the completion notification of the scale-up or the scale-down from the	Upon reception of the completion notification of the scale-up or the scale-down from the second computer of the standby system, the controller in the Exemplary Dell Servers controls to execute system switching of the redundant system to switch the second computer operating as the standby

Claim 1	Exemplary Dell Servers
second computer of the standby system, the controller controls to execute system switching of the redundant system to switch the second computer operating as the standby system undergoing the scale-up or scale-down to a new active system and to switch the first computer operating as the active system to a new standby system.	system undergoing the scale-up or scale-down to a new active system and to switch the first computer operating as the active system to a new standby system. *See, e.g.,*: Primary and Secondary Hosts Add to Library RSS Download PDF Feedback Updated on 07/20/2020 Selected product version: VMware vSphere 7.0 When you add a host to a vSphere HA cluster, an agent is uploaded to the host and configured to communicate with other agents in the cluster. Each host in the cluster functions as a primary host or a secondary host. When vSphere HA is enabled for a cluster, all active hosts (that are not in standby, maintenance mode or not disconnected) participate in an election to choose the cluster's primary host. The host that mounts the greatest number of datastores has an advantage in the election. Only one primary host typically exists per cluster and all other hosts are secondary hosts. If the primary host fails, is shut down or put in standby mode, or is removed from the cluster a new election is held. The primary host in a cluster has several responsibilities: • Monitoring the state of secondary hosts. If a secondary host fails or becomes unreachable, the primary host identifies which virtual machines must be restarted. • Monitoring the power state of all protected virtual machines. If one virtual machine fails, the primary host ensures that it is restarted. Using a local placement engine, the primary host also determines where the restart takes place. • Managing the lists of cluster hosts and protected virtual machines. • Acting as the vCenter Server management interface to the cluster and reporting the cluster health state. The secondary hosts primarily contribute to the cluster by running virtual machines locally, monitoring their runtime states, and reporting state updates to the primary host. A primary host can also run and monitor virtual machines. Both secondary hosts and primary hosts implement the VM and Application Monitoring features. One of the functions performed by the primary host is to orchestrate restarts of protected virtual machines. A virtual machine is protected by a primary host after vCenter Server observes that the virtual machine's power state has changed from powered off to powered on in response to a user action. The primary host persists the list of protected virtual machines in the cluster's datastores. A newly elected primary host uses this information to determine which virtual machines to protect. "Primary and Secondary Hosts."

Claim 1	Exemplary Dell Servers

	Virtual Machine Resiliency
	
	Compute Resiliency Windows Server 2016 includes increased virtual machines compute resiliency to help reduce intra-cluster communication issues in your compute cluster as follows:
	
	• **Resiliency options available for virtual machines:** You can now configure virtual machine resiliency options that define behavior of the virtual machines during transient failures:
	
	  ○ **Resiliency Level:** Helps you define how the transient failures are handled.
	
	  ○ **Resiliency Period:** Helps you define how long all the virtual machines are allowed to run isolated.
	
	• **Quarantine of unhealthy nodes:** Unhealthy nodes are quarantined and are no longer allowed to join the cluster. This prevents flapping nodes from negatively effecting other nodes and the overall cluster.
	
	"What's New in Failover Clustering."

Claim 1	Exemplary Dell Servers
	> Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2 A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service. Failover Clustering has many practical applications, including: • Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines • Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V "Failover Clustering Overview."

Claim 1	Exemplary Dell Servers

Health Service faults

Article • 08/25/2021

> Applies to: Windows Server 2016

What are faults

The Health Service constantly monitors your Storage Spaces Direct cluster to detect problems and generate "faults". One new cmdlet displays any current faults, allowing you to easily verify the health of your deployment without looking at every entity or feature in turn. Faults are designed to be precise, easy to understand, and actionable.

Each fault contains five important fields:

- Severity
- Description of the problem
- Recommended next step(s) to address the problem
- Identifying information for the faulting entity
- Its physical location (if applicable)

For example, here is a typical fault:

```
Severity: MINOR
Reason: Connectivity has been lost to the physical disk.
Recommendation: Check that the physical disk is working and properly
connected.
Part: Manufacturer Contoso, Model XYZ9000, Serial 123456789
Location: Seattle DC, Rack B07, Node 4, Slot 11
```

Claim 1	Exemplary Dell Servers
	"Health Service Faults." High Availability Clusters **High-availability clusters** (also known as **HA clusters** , **failover clusters** or **Metroclusters Active/Active**) are groups of computers that support server applications that can be reliably utilised with a minimum amount of down-time. They operate by using high availability software to harness redundant computers in groups or clusters that provide continued service when system components fail. Without clustering, if a server running a particular application crashes, the application will be unavailable until the crashed server is fixed. HA clustering remedies this situation by detecting hardware/software faults, and immediately restarting the application on another system without requiring administrative intervention, a process known as failover. As part of this process, clustering software may configure the node before starting the application on it. For example, appropriate file systems may need to be imported and mounted, network hardware may have to be configured, and some supporting applications may need to be running as well. HA clusters are often used for critical databases, file sharing on a network, business applications, and customer services such as electronic commerce websites. "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers
	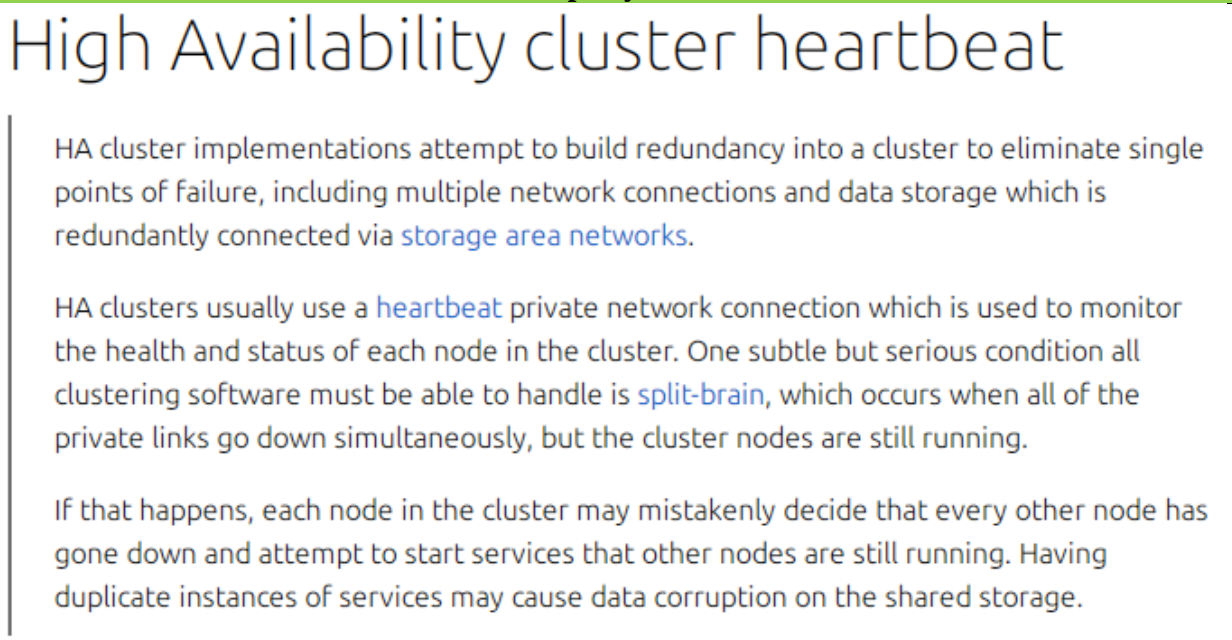 "Ubuntu – Introduction to High Availability."

Claim 1	Exemplary Dell Servers

Linux Clustering Concepts

A cluster is a group of computers (nodes) which work together to provide a shared solution. At a high level, a cluster can be viewed as having three parts (often defined as cluster stack).

Basic concepts

- **Resources:** These are the reason for the cluster's being the services that need to be kept highly available.
- **Resource Agents:** These are scripts operating system components that start, stop, and monitor resources, given a set of resource parameters.
- **Fence Agents:** These are scripts that execute a node fencing actions, given a target and fence device.
- **Fencing:** The ability to disable nodes.
- **Quorum:** Encapsulates the ability to determine whether the cluster can continue to operate safely or not.

Cluster types

The four types are as follows:

- **High Availability (HA):** Used for Fault Tolerance to keep server services available to employees or customers.
- **Load Balancing:** Balances the load between multiple systems when a service needs to be available to numerous systems at once (can be used for other three types of Clusters).
- **Distributed:** Jobs will be managed by different systems.
- **Parallel (Beowulf):** Jobs are managed by multiple processors on multiple systems.

As well by the configuration there are several types of clusters:

- **Manual clustering:** Lets you classify, merge, split clusters manually if the output of the automatic spike sorting algorithms are not satisfactory.
- **Merging clusters:** When multiple clusters seem to correspond to the same unit.
- **Splitting clusters:** You can create a new cluster by drawing a polygon around a set of spikes in the feature view, the amplitude view, the template amplitude view, or the spike attribute views.

"Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	Reasons to have a cluster There are different reasons for using a cluster. We use them to provide a resilient, highly available backend for solutions. This is primarily for backend services like MySQL, NFS or Redis where the service can be behind a front end of web servers. High availability is provided via automatic failover - if there is a failure on a cluster node, the clustered services running on that node will automatically relocate away to a node that is running properly. Clusters will generally use shared storage (SAN) so that data will be persistent when services move between nodes. In the event if a node becomes unresponsive, typically that node will be rebooted (fenced) by other nodes in order to preserve data integrity of the shared storage and floating IP ownership. **Linux Open Source High Availability Clustering** Some Linux operating system vendors offer clustering software, such as SUSE Linux HAE; Red Hat Enterprise Linux (RHEL); and Oracle Real Application Clusters (RAC). While they allow you to create a failover cluster, they present a variety of challenges, thist is highly manual and prone to human error. Linux open-source HA extensions require a high degree of technical skill, creating complexity and reliability issues that challenge most operators. "Rackspace – Clustering concepts."

Claim 1	Exemplary Dell Servers
	## What are high-availability clusters? High-availability architectures run active failover clusters, so there is built-in redundancy and failover and—hopefully—zero downtime. Within the cluster, nodes are monitored not just for availability, but for overall performance of applications, services, and network. Because there is shared storage, there is no data loss if a node goes down, because all cluster nodes work from the same data source. Load balancing can be used to manage traffic for best performance. Outside those broad characteristics, high-availability clusters can be designed for more specialized activities, depending on the priorities and activities within the IT infrastructure. The Red Hat Enterprise Linux High Availability Add-on, for example, has four default configurations: • **High availability**: focuses on uptime and availability • **High performance**: for high speed, concurrent operations • **Load balancing**: for cost-effective scalability • **Storage**: for resilient data management In real-life environments, the high-availability systems would incorporate aspects of those focus elements. "Red Hat – What is high availability."